

Express Mail ET808936103US

**APPLICATION
FOR
UNITED STATES LETTERS PATENT**

APPLICANT NAME: Dennis J. Carroll

**TITLE: ALGORITHM FOR SORTING BIT SEQUENCES IN
LINEAR COMPLEXITY**

DOCKET NO. : END920030031US1

INTERNATIONAL BUSINESS MACHINES CORPORATION

ALGORITHM FOR SORTING BIT SEQUENCES IN LINEAR COMPLEXITY

Background of the Invention

1. Technical Field

The present invention generally relates to an algorithm for sorting bit sequences, and in particular to an algorithm for sorting bit sequences in linear complexity.

2. Related Art

In the current state of the art with respect to sorting words (i.e., integers, strings, etc.), the fastest known algorithms have an execution speed proportional to $N_w \log N_w$ (i.e., of order $N_w \log N_w$), wherein N_w denotes the number of words to be sorted. The well-known Quicksort algorithm is an in-place sort algorithm (i.e., the sorted items occupy the same storage as the original items) that uses a divide and conquer methodology. To solve a problem by divide and conquer on an original instance of a given size, the original instance is divided into two or more smaller instances; each of these smaller instances is recursively solved (i.e., similarly divided), and the resultant solutions are combined to produce a solution for the original instance. To implement divide and conquer, Quicksort picks an element from the array (the pivot), partitions the remaining elements into those greater than and less than this pivot, and recursively sorts the partitions. The execution speed of Quicksort is a function of the sort ordering that is present in the array of words to be sorted. For a totally random distribution of words to be sorted, Quicksort's execution speed is proportional to $N_w \log N_w$. In some cases in which the words to be sorted

deviate from perfect randomness, the execution speed may deteriorates relative to $N_w \log N_w$ and is proportional to $(N_w)^2$ in the worst case.

Given, the enormous execution time devoted to sorting a large number of integers, strings, etc. for extensively used applications such as spreadsheets, database applications, etc.,
5 there is a need for a sort algorithm having an execution speed of order less than $N_w \log N_w$.

Summary of the Invention

The present invention provides a method, computer program product, and associated algorithm for sorting S sequences of binary bits in ascending or descending order of a value associated with each sequence and in a time period denoted as a sorting execution time, said S
10 sequences being stored in a memory device of the computer system prior to said sorting, S being at least 2, each sequence of the S sequences comprising contiguous fields of bits, said sorting comprising executing program code at nodes of a linked execution structure, said executing program code being performed in a sequential order with respect to said nodes, said executing including:

15 masking the contiguous fields of the S sequences in accordance with a mask whose content is keyed to the field being masked, said sequential order being a function of an ordering of masking results of said masking; and

 outputting each sequence of the S sequences or a pointer thereto to an output array of the memory device whenever said masking places said each sequence in a leaf node of the linked
20 execution structure.

The present invention provides a method, computer program product, and associated algorithm for sorting S sequences of binary bits in ascending or descending order of a value associated with each sequence and in a time period denoted as a sorting execution time, said S sequences being stored in a memory device of the computer system prior to said sorting, S being
5 at least 2, each sequence of the S sequences comprising K contiguous fields denoted left to right as F_1, F_2, \dots, F_K with corresponding field widths of W_1, W_2, \dots, W_K , said sorting comprising the steps of:

designating S memory areas of the memory device as A_1, A_2, \dots, A_S ;

setting an output index $P = 0$ and a field index $Q = 0$;

10 providing a node E having S elements stored therein, said S elements consisting of the S sequences or S pointers respectively pointing to the S sequences; and

executing program code, including determining a truth or falsity of an assertion that the elements in node E collectively include or point to no more than one unique sequence U of the S sequences, and if said assertion is determined to be false:

15 then generating C child nodes from node E , each child node including all elements in node E having a unique value of field F_{Q+1} , said child nodes denoted as E_0, E_1, \dots, E_{C-1} having associated field F_{Q+1} values of V_0, V_1, \dots, V_{C-1} , said child nodes E_0, E_1, \dots, E_{C-1} being sequenced such that $V_0 < V_1 < \dots < V_{C-1}$, said generating followed by incrementing Q by 1, said incrementing Q followed by iterating from an index $I=0$ to $I=C-1$ in steps of 1,
20 wherein iteration I includes setting $E=E_I$ followed by executing the program code recursively at a next level of recursion for the node E ;

else for each element in node E: incrementing P by 1, next storing in A_p either U
or the element pointing to U, and lastly if the program code at all of said levels of
recursion has not been fully executed then resuming execution of said program code at
the most previous level of recursion at which the program code was partially but not fully
5 executed else exiting the algorithm.

The present invention provides a method, computer program product, and associated
algorithm for sorting S sequences of binary bits in ascending or descending order of a value
associated with each sequence and in a time period denoted as a sorting execution time, said S
sequences being stored in a memory device of the computer system prior to said sorting, S being
10 at least 2, each sequence of the S sequences comprising K contiguous fields denoted left to right
as F_1, F_2, \dots, F_K with corresponding field widths of W_1, W_2, \dots, W_K , said sorting comprising the
steps of:

designating S memory areas of the memory device as A_1, A_2, \dots, A_S ;

setting an output index $P = 0$ and a field index $Q = 0$;

15 providing a node E having S elements stored therein, said S elements consisting of the S
sequences or S pointers respectively pointing to the S sequences; and

counter-controlled looping through program code, said looping including iteratively
executing said program code within nested loops, said executing said program code including
determining a truth or falsity of an assertion that the elements in node E collectively include or
20 point to no more than one unique sequence U of the S sequences, and if said assertion is

determined to be false:

then generating C child nodes from node E, each child node including all elements in node E having a unique value of field F_{Q+1} , said child nodes denoted as E_0, E_1, \dots, E_{C-1} having associated field F_{Q+1} values of V_0, V_1, \dots, V_{C-1} , said child nodes E_0, E_1, \dots, E_{C-1} being sequenced such that $V_0 < V_1 < \dots < V_{C-1}$; said generating followed by incrementing Q by 1, said incrementing Q followed by iterating from an index $I=0$ to $I=C-1$ in steps of 1, wherein iteration I includes setting $E=E_I$ followed by returning to said counter-controlled looping;

else for each element in node E: incrementing P by 1, next storing in A_P either U or the element pointing to U, and lastly if all iterations of said outermost loop have not been executed then returning to said counter-controlled looping else exiting from said algorithm.

The present invention advantageously provides a sort algorithm having an execution speed of order less than $N_w \log N_w$.

Brief Description of the Drawings

FIG. 1 depicts a path through a linked execution structure, in accordance with embodiments of the present invention.

FIG. 2 depicts paths through a linked execution structure for sorting integers, in accordance with embodiments of the present invention.

FIG. 3 depicts FIG. 2 with the non-existent nodes deleted, in accordance with embodiments of the present invention.

FIG. 4 depicts paths through a linked execution structure for sorting strings with each path terminated at a leaf node, in accordance with embodiments of the present invention.

5 FIG. 5 is a flow chart for linear sorting under recursive execution, in accordance with embodiments of the present invention.

FIG. 6 is a flow chart for linear sorting under counter-controlled looping, in accordance with embodiments of the present invention.

10 FIGS. 7A-7D comprise source code for linear sorting of integers under recursive execution, in accordance with embodiments of the present invention.

FIGS. 8A-8D comprise source code for linear sorting of strings under recursive execution, in accordance with embodiments of the present invention.

FIG. 9 illustrates a computer system for sorting sequences of bits, in accordance with embodiments of the present invention.

15 FIG. 10 is a graph depicting the number of moves used in sorting integers for a values range of 0-9,999,999, using Quicksort and also using the linear sort of the present invention.

FIG. 11 is a graph depicting the number of compares used in sorting integers for a values range of 0-9,999,999, using Quicksort and also using the linear sort of the present invention.

20 FIG. 12 is a graph depicting the number of moves used in sorting integers for a values range of 0-9,999, using Quicksort and also using the linear sort of the present invention.

FIG. 13 is a graph depicting the number of compares used in sorting integers for a values

range of 0-9,999, using Quicksort and also using the linear sort of the present invention.

FIG. 14 is a graph depicting sort time used in sorting integers for a values range of 0-9,999,999, using Quicksort and also using the linear sort of the present invention.

FIG. 15 is a graph depicting sort time used in sorting integers for a values range of 0-9,999, using Quicksort and also using the linear sort of the present invention.

FIG. 16 is a graph depicting memory usage for sorting fixed-length bit sequences representing integers, using Quicksort and also using the linear sort of the present invention.

FIG. 17 is a graph depicting sort time using Quicksort for sorting strings, in accordance with embodiments of the present invention.

FIG. 18 is a graph depicting sort time using a linear sort for sorting strings, in accordance with embodiments of the present invention.

FIGS. 19-24 is a graph depicting sort time used in sorting integers, using Quicksort and also using the linear sort of the present invention, wherein the sort time is depicted as a function of mask width and maximum value that can be sorted.

Detailed Description of the Invention

The detailed description is presented *infra* in three sections. The first section, in conjunction with FIG. 1, comprises an Introduction to the present invention, including assumptions, terminology, features, etc. of the present invention. The second section, in conjunction with FIGS. 2-9 comprises a Sort Algorithm detailed description in accordance with the present invention. The third section, in conjunction with FIGS. 10-24, relates to Timing

Tests, including a description and analysis of execution timing test data for the sort algorithm of the present invention in comparison with Quicksort.

Introduction

5 FIG. 1 depicts a path through linked execution structure, in accordance with embodiments of the present invention. The linked execution structure of FIG. 1 is specific to 12-bit words divided into 4 contiguous fields of 3 bits per field. For example, the example word 100011110110 shown in FIG. 1 is divided into the following 4 fields (from left to right): 100, 011, 110, 110. Each field has 3 bits and therefore has a “width” of 3 bits. The sort algorithm of the present invention will utilize a logical mask whose significant bits (for masking purposes) encompass W bits. Masking a sequence of bits is defined herein as extracting (or pointing to) a subset of the bits of the sequence. Thus, the mask may include a contiguous group of ones (i.e., 11 ... 1) and the remaining bits of the mask are each 0; the significant bits of the mask consist of the contiguous group of ones, and the width W of the mask is defined as the number of the significant bits in the mask. Thus, W is referred to as a “mask width”, and the mask width W determines the division into contiguous fields of each word to be sorted. Generally, if the word to be sorted has N bits and if the mask width is W, then each word to be sorted is divided into L fields (or “levels”) such that $L=N/W$ if N is an integral multiple of W, under the assumption that the mask width W is constant. If N is not an integral multiple of W, then the mask width cannot be constant. For example if $N=12$ and $W=5$, then the words to be sorted may be divided into, *inter alia*, 3 fields, wherein going from left to right the three fields have 5 bits, 5 bits, and 2 bits.

In this example, L may be calculated via $L = \text{ceiling}(N/W)$, wherein $\text{ceiling}(x)$ is defined as the smallest integer greater than or equal to x. Thus, the scope of present invention includes an embodiment in which W is a constant width with respect to the contiguous fields of each word to be sorted. Alternatively, the scope of present invention also includes an embodiment in which W is a variable width with respect to the contiguous fields of each word to be sorted. Each word to be sorted may be characterized by the same mask and associated mask width W, regardless of whether W is constant or variable with respect to the contiguous fields.

Although the scope of the present invention permits a variable mask width W as in the preceding example, the example of FIG. 1 as well as the examples of FIGS. 2-4 discussed *infra* use a constant mask width for simplicity. For the example of FIG. 1, $N=12$, $W=3$, and $L=4$. It should be noted that the maximum numerical value that the N bits could have is 2^N-1 . Thus, the maximum value that a 12-bit word could have is 4095.

In FIG. 1, the linked execution structure has a root, levels, and nodes. Assuming a constant mask of width W, the root in FIG. 1 is represented as a generic field of W bits having the form xxx where x is 0 or 1. Thus, the width W of the mask used for sorting is the number of bits (3) in the root. The generic nodes corresponding to the root encompass all possible values derived from the root. Hence the generic nodes shown in FIG. 1 are 000, 001, 010, 011, 100, 101, 110, and 111. The number of such generic nodes is 2^W , or 8 if $W=3$ as in FIG. 1. There are L levels (or “depths”) such that each field of a word corresponds to a level of the linked execution structure. In FIG. 1, the 4 levels (i.e., $L=4$) are denoted as Level 1, Level 2, Level 3, and Level 4.

Consider the example word 100011110110 shown in FIG. 1. Below the root are 8 generic nodes of Level 1, called “child nodes” of the root. The first field of the example word is 100 corresponding to the 100 node in Level 1. Below the 100 node of Level 1 are the 8 generic nodes of Level 2, namely the child nodes of the 100 node of Level 1. The second field of the example word is 011 corresponding to the 011 node in Level 2. Below the 011 node of Level 2 are its 8 child nodes in Level 3. The third field of the example word is 110 corresponding to the 110 node in Level 3. Below the 110 node of Level 3 are its 8 child nodes in Level 4. The fourth field of the example word is 110 corresponding to the 110 node in Level 4. Thus, the path through the linked execution structure for the example word 100011110110 consists of the 100 node of level 1, the 011 child node of Level 2, the 110 child node of Level 3, and the 110 child node of Level 4.

Although not shown in FIG. 1, each node of the linked execution structure at level i potentially has the 2^W child nodes below it at level $i+1$. For example the 000 node at Level 1 has 8 child nodes below it, and each such child nodes has 8 child nodes, etc. Thus the maximum number of nodes of the linked execution structure is $2^W + 2^{2W} + 2^{3W} + \dots + 2^{LW}$, or $(2^{(L+1)W} - 2^W)/(2^W - 1)$. In FIG. 1, the total number of nodes is 4680 for $W=3$ and $L=4$. Since it is not practical to show all nodes of the linked execution structure, FIG. 1 shows only those nodes and their children which illustrate the path of the example word.

The actual nodes of a linked execution structure relative to a group of words to be sorted comprise actual nodes and non-existent nodes. The paths of the words to be sorted define the actual nodes, and the remaining nodes define the non-existent nodes. Thus in FIG. 1, the actual

nodes include 100 node of level 1, the 011 child node of Level 2, the 110 child node of Level 3, and the 110 child node of Level 4. Any other word having a path through the linked execution structure of FIG. 1 defines additional actual nodes.

Another concept of importance is a "leaf node" of the linked execution structure, which is an actual node that is also a terminal node of a path through the linked execution structure. A leaf node has no children. In FIG. 1, 110 node in Level 4 is a leaf node. In the context of the sort algorithm of the present invention, it is also possible to a leaf node at a level other than the deepest Level L. Multiple numbers to be sorted may give rise to a given node having more than one child (i.e., the paths of different numbers to be sorted may intersect in one or more nodes). If a given node of the linked execution structure holds more than one unique word to be sorted, then the algorithm must process the child nodes of the given node. If, however, the given node of the linked execution structure holds no more than one unique word to be sorted, then the given node is a leaf node and the sort algorithm terminates the path at the given node without need to consider the child (if any) of the given node. In this situation, the given node is considered to be a leaf node and is considered to effectively have no children. Thus, it is possible for a leaf node to exist at a level L_1 wherein $L_1 < L$. The concept of such leaf nodes will be illustrated by the examples depicted in FIGS. 2-4, discussed *infra*.

The sort algorithm of the present invention has an execution time that is proportional to $N*Z$, wherein Z is a positive real number such that $1 \leq Z \leq L$. As stated *supra*, N is defined as the number of bits in each word to be sorted, assuming that N is a constant and characterizes each word to be sorted, wherein said assumption holds for the case of an integer sort, a floating point

sort, or a string sort such that the string length is constant. Z is a function of the distribution of leaf nodes in the linked execution structure. The best case of $Z=1$ occurs if all leaf nodes are at level 1. The worst case of $Z=L$ occurs if all leaf nodes occur at Level L . Thus, the execution time for the worst case is proportional to $N*L$, and is thus linear in N with L being a constant that is controlled by a choice of mask width W . Therefore, the sort algorithm of the present invention is designated herein as a “linear sort”. The term “linear sort” is used herein to refer to the sorting algorithm of the present invention.

If the words to be sorted are strings characterized by a variable string length, then the execution time is proportional to $\sum_j W_j N_j$, where N_j is a string length in bits or bytes (assuming that the number of bits per byte is a constant), wherein W_j is a weighting factor that is proportional to the number of strings to be sorted having a string length N_j . The summation \sum_j is from $j=1$ to $j=J$ such that J is the number of unique string lengths in the strings to be sorted. For example consider 60 strings to be sorted such that 30 strings have 3 bytes each, 18 strings have 4 bytes each, and 12 strings have 5 bytes each. For this example, $J=3$, $N_1=3$ bytes, $W_1 \propto 30$, $N_2=4$ bytes, $W_2 \propto 18$, $N_3=5$ bytes, $W_3 \propto 12$ bytes, wherein the symbol “ \propto ” stands for “proportional to”. Thus, the sort execution time is a linear combination of the string lengths N_j (expressed in bits or bytes) of the variable-length strings to be sorted. Accordingly, the sort algorithm of the present invention is properly designated herein as a “linear sort” for the case of sorting variable-length strings.

In light of the preceding discussion, the sort algorithm of the present invention is designated herein as having a sorting execution time for sorting words (or sequences of bits),

wherein said sorting execution time is a linear function of the word length (or sequence length) of the words (or sequences) to be sorted. The word length (or sequence length) may be a constant length expressed as a number of bits or bytes (e.g., for integer sorts, floating point sorts, or string sorts such that the string length is constant). Thus for the constant word length (or
5 sequence length) case, an assertion herein and in the claims that the sorting execution time function is a linear function of the word length (or sequence length) of the words (or sequences) to be sorted means that the sorting execution time is linearly proportional to the constant word length (or sequence length).

Alternatively, the word length (or sequence length) may be a variable length expressed as
10 numbers of bits or bytes (e.g., for string sorts such that the string length is variable). Thus for the constant word length (or sequence length) case, an assertion herein and in the claims that the sorting execution time function is a linear function of the word length (or sequence length) of the words (or sequences) to be sorted means that the sorting execution time is proportional to a linear combination of the unique non-zero values of string length (i.e., $N_j \neq 0$) which characterize the
15 strings to be sorted.

Note that the sorting execution time of the present invention is also a linear (or less than linear) function of S wherein S is the number of sequences to be sorted, as will be discussed *infra*.

Also note that an analysis of the efficiency of the sorting algorithm of the present
20 invention may be expressed in terms of an “algorithmic complexity” instead of in terms of a sorting execution time, inasmuch as the efficiency can be analyzed in terms of parameters which

the sorting execution time depends on such as number of moves, number of compares, etc. This will be illustrated *infra* in conjunction with FIGS. 10-13.

As stated *supra*, $L=N/W$ (if W is constant) and the upper-limiting value V_{UPPER} that may potentially be sorted is 2^N-1 . Consequently, $L = (\log_2 V_{UPPER} + 1)/W$. Interestingly, L is thus
5 dependent upon both W and V_{UPPER} and does not depend on the number of values to be sorted, which additionally reduces the sort execution time. Inspection of the sort algorithm shows that a larger mask width W indicates a less efficient use of memory but provides a faster sort except at the very highest values of W (see FIGS. 19-24 and description thereof). Since the sort execution time depends on W through the dependence of L or Z on W , one can increase the sort execution
10 speed by adjusting W upward in recognition of the fact that a practical upper limit to W may be dictated by memory storage constraints, as will be discussed *infra*.

The sort algorithm of the present invention assumes that: 1) for any two adjacent bits in the value to be sorted, the bit to the left represents a larger magnitude effect on the value than the bit to the right; or 2) for any two adjacent bits in the value to be sorted, the bit to the right
15 represents a larger magnitude effect on the value than the bit to the left. The preceding assumptions permit the sort algorithm of the present invention to be generally applicable to integer sorts and string sorts. The sort algorithm is also applicable to floating point sorts in which the floating point representation conforms to the commonly used format having a sign bit denoting the sign of the floating point number, an exponent field (wherein positive and negative
20 exponents may be differentiated by addition of a bias for negative exponents as will be illustrated *infra*), and a mantissa field, ordered contiguously from left to right in each word to be sorted. The

sort algorithm is also applicable to other data types such as: other floating point representations consistent with 1) and 2) above; string storage such that leftmost bytes represent the length of the string; little endian storage; etc.

The sort algorithm of the present invention includes the following characteristics: 1) the
5 sort execution time varies linearly with N as discussed *supra*; 2) the sort execution time varies
linearly (or less than linearly) with S as discussed *supra*; 3) the values to be sorted are not
compared with one another as to their relative values or magnitudes; 4) the sort execution speed
is essentially independent of the data ordering characteristics (with respect to data value or
magnitude) in the array of data to be sorted; 5) the sort efficiency (i.e., with respect to execution
10 speed) varies with mask width and the sort efficiency can be optimized through an appropriate
choice of mask width; 6) for a given mask width, sort efficiency improves as the data density
increases, wherein the data density is measured by $S/(V_{MAX}-V_{MIN})$, wherein S denotes the number
of values to be sorted, and wherein V_{MAX} and V_{MIN} are, respectively, the maximum and minimum
values within the data to be sorted, so that the sort execution time may vary less than linearly with
15 S (i.e., the sort execution time may vary as S^Y such that $Y < 1$); and 7) although the linked
execution structure of FIG. 1 underlies the methodology of the sort algorithm, the linked
execution structure is not stored in memory during execution of the sort (i.e., only small portions
of the linked execution structure are stored in memory at any point during execution of the sort).

The linked execution structure of the present invention includes nodes which are linked
20 together in a manner that dictates a sequential order of execution of program code with respect to
the nodes. Thus, the linked execution structure of the present invention may be viewed a

program code execution space, and the nodes of the linked execution structure may be viewed as points in the program code execution space. As will be seen in the examples of FIGS. 2-4 and the flow charts of FIGS. 5-6, described *infra*, the sequential order of execution of the program code with respect to the nodes is a function of an ordering of masking results derived from a
5 masking of the fields of the words (i.e., sequences of bits) to be sorted.

The Sort Algorithm

FIG. 2 depicts paths through a linked execution structure for sorting integers, in accordance with embodiments of the present invention. FIG. 2 illustrates a sorting method, using a 2-bit mask, for the eight integers (i.e., $S=8$) initially sequenced in decimal as 12, 47, 44, 37, 03,
10 14, 31, and 44. The binary equivalents of the words to be sorted are shown. Each word to be sorted has 6 bits identified from right to left as bit positions 0, 1, 2, 3, 4, and 5. For this example: $S=8$, $N=6$, $W=2$, and $L=3$. The root is represented as a generic field of $W=2$ bits having the form xx where x is 0 or 1. The generic nodes corresponding to the root are 00, 01, 10, and 11. The number of such generic nodes is 2^W , or 4 for $W=2$ as in FIG. 2. There are 3 levels such that each
15 field of a word to be sorted corresponds to a level of the linked execution structure. In FIG. 2, the 3 levels (i.e., $L=3$) are denoted as Level 1, Level 2, and Level 3. A mask of 110000 is used for Level 1, a mask of 001100 is used for Level 2, and a mask of 000011 is used for Level 3.

The Key indicates that a count of the number of values in each node is indicated with a left and right parenthesis (), with the exception of the root which indicates the form xx of the
20 root. For example, the 00 node of level one has three values having the 00 bits in bit positions 4

and 5, namely the values 12 (001100), 03 (000011), and 14 (001110). The Key also differentiates between actual nodes and non-existent nodes. For example, the actual 01 node in Level 1 is a leaf node containing the value 31, so that the nodes in Levels 2 and 3 that are linked to the leaf node 01 in Level 1 are non-existent nodes which are present in FIG. 2 but could have been omitted from FIG. 2. Note that non-existent nodes not linked to any path are omitted entirely from FIG. 2. For example, the non-existent 11 node in Level 1 has been omitted, since none of the words to be sorted has 11 in bit positions 4 and 5. FIG. 3 depicts FIG. 2 with all non-existent nodes deleted.

The integer sort algorithm, which has been coded in the C-programming language as shown in FIG. 7, is applied to the example of FIG. 2 as follows. An output array A(1), A(2), ..., A(S) has been reserved to hold the outputted sorted values. For simplicity of illustration, the discussion *infra* describes the sort process as distributing the values to be sorted in the various nodes. However, the scope of the present invention includes the alternative of placing pointers to values to be sorted (e.g., in the form of linked lists), instead of the values themselves, in the various nodes. Similarly, the output array A(1), A(2), ..., A(S) may hold the sorted values or pointers to the sorted values.

The mask at each level is applied to a node in the previous level, wherein the root may be viewed as a root level which precedes Level 1, and wherein the root or root level may be viewed as holding the S values to be sorted. In FIG. 2 and viewing the root as holding all eight values to be sorted, the Level 1 mask of 110000 is applied to all eight values to be sorted to distribute the values in the 4 nodes (00, 01, 10, 11) in Level 1 (i.e., based on the bit positions 4 and 5 in the

words to be sorted). The generic nodes 00, 01, 10, 11 are ordered in ascending value (i.e., 0, 1, 2, 3) from left to right at each of Levels 1, 2 and 3, which is necessary for having the sorted values automatically appear outputted sequentially in ascending order of value. It is also necessary to have the 11 bits in the mask shifted from left to right as the processing moves down in level from Level 1 to Level 2 to Level 3, which is why the 11 bits are in bit positions 4-5 in Level 1, in bit positions 2-3 in Level 2, and in bit positions 0-1 in Level 3. Applying the mask (denoted as “MASK”) to a word (“WORD”) means performing the logical operation MASK AND WORD to isolate all words having bits corresponding to “11” in MASK. As shown for Level 1, the 00 node has 3 values (12, 03, 14), the 01 node has 1 value (31), the 10 node has 4 values (47, 44, 37, 44), and the 11 node has zero values as indicated by the absence of the 11 node at Level 1 in FIG. 2. Note that the 10 node in Level 1 has duplicate values of 44. Next, the actual nodes 00, 01, and 10 in Level 1 are processed from left to right.

Processing the 00 node of Level 1 comprises distributing the values 12, 03, and 14 from the 00 node of Level 1 into its child nodes 00, 01, 10, 11 in Level 2, based on applying the Level 2 mask of 001100 to each of the values 12, 03, and 14. Note that the order in which the values 12, 03, and 14 are masked is arbitrary. However, it is important to track the left-to-right ordering of the generic 00, 01, 10, and 11 nodes as explained *supra*. FIG. 2 shows that the 00 node of Level 2 (as linked to the 00 node of Level 1) is a leaf node, since the 00 node of Level 2 has only 1 value, namely 03. Thus, the value 03 is the first sorted value and is placed in the output array element A(1). Accordingly, the 00, 01, 10, and 11 nodes of Level 3 (which are linked to the 00 node of Level 2 which is linked to the 00 node of Level 1) are non-existent nodes. FIG. 2 also

shows that the 11 node of level 2 (as linked to the 00 node of Level 1) has the two values of 12 and 14. Therefore, the values 12 and 14 in the 11 node of level 2 (as linked to the 00 node of Level 1) are to be next distributed into its child nodes 00, 01, 10, 11 of Level 3, applying the Level 3 mask 000011 to the values 12 and 14. As a result, the values 12 and 14 are distributed into the leaf nodes 00 and 10, respectively, in Level 3. Processing in the order 00, 01, 10, 11 from left to right, the value 12 is outputted to A(2) and the value 14 is outputted to A(3).

FIG. 2 shows that the 01 node of Level 1 is a leaf node, since 31 is the only value contained in the 01 node of Level 1. Thus, the value of 31 is outputted to A(4). Accordingly, all nodes in Level 2 and 3 which are linked to the 01 node of Level 1 are non-existent nodes.

Processing the 10 node of Level 1 comprises distributing the four values 47, 44, 37, and 44 from the 10 node of Level 1 into its child nodes 00, 01, 10, 11 in Level 2, based on applying the Level 2 mask of 001100 to each of the values 47, 44, 37, and 44. FIG. 2 shows that the 01 node of Level 2 (as linked to the 10 node of Level 1) is a leaf node, since the 01 node of Level 2 has only 1 value, namely 37. Thus, the value 37 is placed in the output array element A(5).

Accordingly, the 00, 01, 10, and 11 nodes of Level 3 which are linked to the 01 node of Level 2 which is linked to the 10 node of Level 1 are non-existent nodes. FIG. 2 also shows that the 11 node of level 2 (as linked to the 10 node of Level 1) has the three values of 47, 44, and 44. Therefore, the values 47, 44, and 44 in the 11 node of level 2 (as linked to the 10 node of Level 1) are to be next distributed into its child nodes 00, 01, 10, 11 of Level 3 (from left to right), applying the Level 3 mask 000011 to the values 47, 44, and 44. As a result, the duplicate values of 44 and 44 are distributed into the leaf nodes 00 in Level 3, and the value of 47 is distributed

into the leaf node 11 in level 3. Processing in the order 00, 01, 10, 11 from left to right, the value 44 is outputted to A(6), the duplicate value 44 is outputted to A(7), and the value 47 is outputted to A(8). Thus, the output array now contains the sorted values in ascending order or pointers to the sorted values in ascending order, and the sorting has been completed.

5 While the preceding discussion of the example of FIG. 2 considered the words to be sorted to be integers, each of the words to be sorted could be more generally interpreted as a contiguous sequence of binary bits. The sequence of bits could be interpreted as an integer as was done in the discussion of FIG. 2 *supra*. The sequence of bits could alternatively be interpreted as a character string, and an example of such a character string interpretation will be
10 discussed *infra* in conjunction with FIG. 4. Additionally, the sequence could have been interpreted as a floating point number if the sequence had more bits (i.e., if N were large enough to encompass a sign bit denoting the sign of the floating point number, an exponent field, and a mantissa field). Thus, the sorting algorithm is generally an algorithm for sorting sequences of bits whose interpretation conforms to the assumptions stated *supra*. It should be noted, however,
15 that if the sequences are interpreted as numbers (i.e., as integers or floating point numbers) then the word length (in bits) N must be constant. If the sequences are interpreted as character strings, however, then the word length N is not required to be constant and the character strings to be sorted may have a variable length.

 An important aspect of the preceding sort process is that no comparisons were made
20 between the values to be sorted, which has the consequence of saving an enormous amount of processing time that would otherwise have been expended had such comparisons been made.

The sort algorithm of the present invention accomplishes the sorting in the absence of such comparisons by the masking process characterized by the shifting of the 11 bits as the processing moves down in level from Level 1 to Level 2 to Level 3, together with the left to right ordering of the processing of the generic 00, 01, 10, 11 nodes at each level. The fact that the output array
5 A(1), A(2), ..., A(8) contains sorted values in ascending order is a consequence of the first assumption that for any two adjacent bits in the value to be sorted, the bit to the left represents a larger magnitude effect on the value than the bit to the right. If the alternative assumption had been operative (i.e., for any two adjacent bits in the value to be sorted, the bit to the right represents a larger magnitude effect on the value than the bit to the left), then the output array
10 A(1), A(2), ..., A(8) would contain the same values as under the first assumption; however the sorted values in A(1), A(2), ..., A(8) would be in descending order.

The preceding processes could be inverted and the sorted results would not change except possibly the ascending/descending aspect of the sorted values in A(1), A(2), ..., (8). Under the inversion, the generic bits would be processed from right to left in the ordered sequence: 00, 01, 10,
15 11 (which is equivalent to processing the ordered sequence 11, 10, 01, 00 from left to right). As a result, the output array A(1), A(2), ..., A(8) would contain sorted values in descending order as a consequence of the first assumption that for any two adjacent bits in the value to be sorted, the bit to the left represents a larger magnitude effect on the value than the bit to the right. However under the inversion and if the alternative assumption had been operative (i.e., for any two
20 adjacent bits in the value to be sorted, the bit to the right represents a larger magnitude effect on the value than the bit to the left), then the output array A(1), A(2), ..., A(8) would contain the

sorted values in ascending order.

The preceding process assumed that the mask width W is constant. For example, $W=2$ for the example of FIG. 2. However, the mask width could be variable (i.e., as a function of level or depth). For example consider a sort of 16 bit words having mask widths of 3, 5, 4, 4 at levels

5 1, 2, 3, 4, respectively. That is, the mask at levels 1, 2, 3, and 4 may be, *inter alia*,

1110000000000000, 0001111100000000, 0000000011110000, and 0000000000001111,

respectively. Generally, for N -bit words to be sorted and L levels of depth, the mask widths W_1 ,

W_2 , ..., W_L corresponding to levels 1, 2, ..., L , respectively, must satisfy: $W_1 + W_2 + \dots + W_L \leq N$.

It is always possible have masks such that $W_1 + W_2 + \dots + W_L = N$. However, an improvement in

10 efficiency may be achieved for the special case in which all numbers to be sorted have 0 in one or more contiguous leftmost bits, as will be illustrated *infra*. In said special case, said leftmost bits having 0 in all words to be sorted would not be masked and consequently $W_1 + W_2 + \dots + W_L < N$.

There are several reasons for having a variable mask width. A first reason for having a
15 variable mask width W is that it may not be logically possible to have a constant mask width if $L > 1$, such as for the case of N being a prime number. For example, if $N=13$, then there does not exist an integer L of at least 2 such that N/L is an integer. In theory, it is potentially possible to choose $W=N$ even if N is a prime number. However, memory constraints may render the choice of $W=N$ unrealistic as will be discussed next.

20 A second reason for having a variable mask width W , even if it logically possible for W to be constant with $L > 1$, is that having a variable W may reduce the sort execution time

inasmuch as the sort execution time is a function of W as stated *supra*. As W is increased, the number of levels may decrease and the number of nodes to be processed may likewise decrease, resulting in a reduction of processing time. However, the case of sufficiently large W may be characterized by a smallest sort execution time, but may also be characterized by prohibitive memory storage requirements and may be impracticable (see *infra* FIG. 16 and discussion thereof). Thus in practice, it is likely that W can be increased up to a maximum value above which memory constraints become controlling. Thus the case of $L > 1$ is highly likely, and two or more mask widths will exist corresponding to two or more levels. As will be seen from the analysis of timing test data discussed in conjunction with FIGS. 19-24 discussed *infra*, the sort efficiency with respect to execution speed is a function not only of mask width but also of the data density as measured by $S/(V_{MAX}-V_{MIN})$. Moreover, the mask width and the data density do not independently impact the sort execution speed. Instead the mask width and the data density are coupled in the manner in which they impact the sort execution speed. Therefore, it may be possible to fine tune the mask width as a function of level in accordance with the characteristics (e.g., the data density) of the data to be sorted.

Another improvement in sort execution timing may result from finding the highest or maximum value V_{MAX} to be sorted and then determine if V_{MAX} is of such a magnitude that N can be effectively reduced. For example, if 8-bit words are to be sorted and V_{MAX} is determined to have the value 00110101, then bits 7-8 of all words to be sorted have 00 in the leftmost bits 6-7. Therefore, bits 7-8 do not have to be processed in the sorting procedure. To accomplish this, a mask could be employed in a three-level sorting scheme having $N=8$, $L=3$, $W_1=2$, $W_2=2$ and

$W_3=2$. The masks for this sorting scheme are 00110000 for level 1, 00001100 for level 2, and 00000011 for level 3. Although $N=8$ technically prevails, the actual sort time will be reflective of $N=6$ rather than $N=8$, because the masks prevent bits 6-7 from being processed.

Similarly, one could find a lowest or minimum value V_{MIN} to be sorted and then
5 determine if V_{MIN} is of such a magnitude that N can be effectively reduced. For example, if 8-bit words are to be sorted and V_{MIN} is determined to have the value 10110100, then bits 0-1 of all words to be sorted have 00 in the rightmost bits 0-1. Therefore, bits 0-1 do not have to be processed in the sorting procedure. To accomplish this, a variable width mask could be employed in a three-level sorting scheme having $N=8$, $L=3$, $W_1=2$, $W_2=2$ and $W_3=2$. The masks
10 for this sorting scheme are 11000000 for level 1, 00110000 for level 2, and 00001100 for level 3. Although $N=8$ technically in this scheme, the actual sort time will be reflective of $N=6$ rather than $N=8$, because the masks prevent bits 0-1 from being processed at all.

Of course, it may be possible to utilize both V_{MAX} and V_{MIN} in the sorting to reduce the effective value of N . For example, if 8-bit words are to be sorted and V_{MAX} is determined to have
15 the value 00110100 and V_{MIN} is determined to have the value 00000100, then bits 7-8 of all words to be sorted have 00 in the leftmost bits 6-7 and bits 0-1 of all words to be sorted have 00 in the rightmost bits 0-1. Therefore, bits 7-8 and 0-1 do not have to be processed in the sorting procedure. To accomplish this, a constant width mask could be employed in a two-level sorting scheme having $N=8$, $L=2$, and $W=2$. The masks for this sorting scheme are 00110000 for level 1
20 and 00001100 for level 2. Although $N=8$ technically in this scheme, the actual sort time will be reflective of $N=4$ rather than $N=8$, because the masks prevent bits 6-7 and 0-1 from being

processed at all.

The integer sorting algorithm described *supra* in terms of the example of FIG. 2 applies generally to integers. If the integers to be sorted are all non-negative, or are all negative, then the output array A(1), A(2), ..., will store the sorted values (or pointers thereto) as previously

5 described. However, if the values to be sorted are in a standard signed integer format with the negative integers being represented as a two's complement of the corresponding positive integer, and if the integers to be sorted include both negative and non-negative values, then output array A(1),A(2), ... stores the negative sorted integers to the right of the non-negative sorted integers..

For example the sorted results in the array A(1), A(2), ... may appear as: 0, 2, 5, 8, 9, -6, -4, -2,

10 and the algorithm could test for this possibility and reorder the sorted results as: -6, -4, -2, 0, 2, 5, 8, 9.

The sorting algorithm described *supra* will correctly sort a set of floating point numbers in which the floating point representation conforms to the commonly used format having a sign bit, an exponent field, and a mantissa field ordered contiguously from left to right in each word to
15 be sorted. The standard IEEE 754 format represents a single-precision real number in the following 32-bit floating point format:

Sign Bit (1 bit)	Exponent Field (8 bits)	Mantissa Field (23 bits)
------------------	-------------------------	--------------------------

IEEE 754 requires the exponent field to have a +127 (i.e., 01111111) bias for positive exponents and no bias for negative exponents. The exponent field bits satisfy the previously stated

assumption that for any two adjacent bits in the value to be sorted, the bit to the left represents a larger magnitude effect on the value than the bit to the right, as may be seen in the following table for the exponents of -2, -1, 0, +1, and +2.

Exponent Value	Exponent Field Bits
-2	01111101
-1	01111110
0	01111111
1	10000000
2	10000001

The number of bits in the exponent and mantissa fields in the above example is merely illustrative. For example, the IEEE 754 representation of a double-precision floating point number has 64 bits (a sign bit, an 11-bit exponent, and a 52-bit mantissa) subject to an exponent bias of +1023. Generally, the exponent and mantissa fields may each have any finite number of bits compatible with the computer/processor hardware being used and consistent with the degree of precision desired. Although the sign bit is conventionally 1 bit, the sort algorithm of the present invention will work correctly even if more than one bit is used to describe the sign. It is assumed herein that the position of the decimal point is in a fixed position with respect to the bits of the mantissa field and the magnitude of the word is modulated by the exponent value in the exponent field, relative to the fixed position of the decimal point. As illustrated *supra*, the exponent value may be positive or negative which has the effect of shifting the decimal point to

the left or to the right, respectively.

Due to the manner in which the sign bit and exponent field affect the value of the floating-point word, a mask may be used to define a field that includes any contiguous sequence of bits. For example, the mask may include the sign bit and a portion of the exponent field, or a portion of the exponent field and a portion of the mantissa field, etc. In the 32-bit example *supra*, for example, the sorting configuration could have 4 levels with a constant mask width of 8 bits: $N=32$, $L=4$, and $W=8$. The mask for level 1 is 11111110_{24} , wherein 0_{24} represents 24 consecutive zeroes. The mask for level 2 is 0000000011111110_{16} , wherein 0_{16} represents 16 consecutive zeroes. The mask for level 3 is $0_{16}1111111100000000$. The mask for level 4 is $0_{24}11111111$. Thus the mask for level 1 includes the sign bit and the 7 leftmost bits of the exponent field, the mask at level 2 includes the rightmost bit of the exponent field and the 7 leftmost bits of the mantissa field, and the mask for levels 3 and 4 each include 8 bits of the mantissa field.

If the floating point numbers to be sorted include a mixture of positive and negative values, then the sorted array of values will have the negative sorted values to the right of the positive sorted values in the same hierarchical arrangement as occurs for sorting a mixture of positive and negative integers described *supra*.

FIG. 4 depicts paths through a linked execution structure for sorting strings with each path terminated at a leaf node, in accordance with embodiments of the present invention. In FIG. 4, thirteen strings of 3 bytes each are sorted. The 13 strings to be sorted are: 512, 123, 589, 014, 512, 043, 173, 179, 577, 152, 256, 167, and 561. Each string comprises 3 characters selected

from the following list of characters: 0, 1, 2, 3, 4, 5, 6, 7, 8, and 9. Each character consists of a byte, namely 8 bits. Although in the example of FIG. 4 a byte consists of 8 bits, a byte may generally consist of any specified number of bits. The number of potential children (i.e., child nodes) at each node is 2^b where b is the number of bits per byte. Thus in FIG. 4, each node potentially has 256 (i.e., 2^8) children. The sequence 014, 043, 123, ... at the bottom of FIG. 4 denoted the strings in their sorted order.

In FIG. 4, the string length is constant, namely 3 characters or 24 bits. Generally, however, the string length may be variable. The character string defines a number of levels of the linked execution structure that is equal to the string length as measured in bytes. There is a one-to-one correspondence between byte number and level number. For example, counting left to right, the first byte corresponds to level 1, the second byte corresponds to level 2, etc. Thus, if the string length is variable then the maximum number of levels L of the linked execution structure is equal to the length of the longest string to be sorted, and the processing of any string to be sorted having a length less than the maximum level L will reach a leaf node at a level less than L .

The mask width is a constant that includes one byte, and the boundary between masks of successive levels coincide with byte boundaries. Although the sorting algorithm described in conjunction with the integer example of FIG. 2 could be used to sort the character strings of FIG. 4, the sorting algorithm to sort strings could be simplified to take advantage of the fact that mask boundaries coincide with byte boundaries. Rather than using an explicit masking strategy, each individual byte may be mapped into a linked list at the byte's respective level within the linked

execution structure. Under this scheme, when the processing of a string reaches a node corresponding to the rightmost byte of the string, the string has reached a leaf node and can then be outputted into the sorted list of strings. For example, a programming language with uses length/value pairs internally for string storage can compare the level reached with the string's length (in bytes) to determine when that the string has reached a leaf node. The preceding scheme is an implicit masking scheme in which the mask width is equal to the number of bits in a character byte. Alternatively, the algorithm could use an explicit masking scheme in which any desired masking configuration could be used (e.g., a mask could encompass bits of two or more bytes). Thus, a masking strategy is always being used, either explicitly or implicitly.

In FIG. 4, the sorting of the thirteen strings 3-byte strings are characterized by $S=13$, $N=24$ (i.e. 3 bytes x 8bits/byte), $W=8$ (i.e., 1 byte), and $L=3$. Shown in each node is a mask associated with the node, and the strings whose path passes through the node. The mask in each node is represented as a sequence of bytes and each byte might may be one of the following three unique symbols: X, x, and h where h represents one of the characters 0, 1, 2, 3, 4, 5, 6, 7, 8, 9..

The position within the mask of the symbol X is indicative of the location (and associated level) of child nodes next processed. The X is used to mask various strings, as will be described *infra*, by setting X equal to the mask character; thus if X is being used to isolate strings having "5" in the masked position of the strings then $X="5"$ will characterize the mask. The symbol "h" and its position in the mask indicates that the strings in the node each have the character represented by "h" in the associated position. The position within the mask of the symbol "x" indicates the location (and associated level) of the mask representative of other child nodes (e.g.,

“grandchildren”) to be subsequently processed.

The strings shown in each node in FIG. 4 each have the form H: s(1), s(2), ..., wherein H represents a character of the string in the byte position occupied by X, and wherein s(1), s(2), ... are strings having the character represented by H in the byte position occupied by X. For example, in the node whose mask is 0Xx, the string denoted by 1:014 has “0” in byte position 1 and “1” in byte position 2, and the string denoted by 4:043 has “0” in byte position 1 and “4” in byte position 2. As another example, in the node whose mask is 17X, the string denoted by 3:173 has “1” in byte position 1, “7” in byte position 2, and “3” in byte position 3, whereas the string denoted by 9:179 has “1” in byte position 1, “7” in byte position 2, and “9” in byte position 3.

The method of sorting the strings of FIG. 4 follows substantially the same procedure as was described *supra* for sorting the integers of FIG. 2. The string sort algorithm, which has been coded in the C-programming language as shown in FIG. 8, is applied to the example of FIG. 4 as follows. Similar to FIG. 2, an output array A(1), A(2), ..., A(S) has been reserved to hold the outputted sorted values. For simplicity of illustration, the discussion *infra* describes the sort process as distributing the values to be sorted in the various nodes. However, the scope of the present invention includes the alternative of placing pointers to values to be sorted (e.g., in the form of linked lists), instead of the values themselves, in the various nodes. Similarly, the output array A(1), A(2), ..., A(S) may hold the sorted values or pointers to the sorted values.

First, the root node mask of Xxx is applied to all thirteen strings to be sorted to distribute the strings in the 10 nodes 0Xx, 1Xx, ..., 9Xx, resulting of the extraction and storage of the

strings to be sorted and their identification with the first byte of 0, 1, 2, 3, 4, 5, 6, 7, 8, or 9.

Applying the mask a string may be accomplished by ANDing the mask with the string to isolate the strings having a byte corresponding to the byte position of X in the mask to identify the child

nodes. As another approach, the character bytes of a string could be pointed to or extracted from

5 the string by use of a string array subscript, wherein the string array subscript serves as the mask by providing the functionality of the mask. Masking a sequence of bits is defined herein as

extracting (or pointing to) a subset of the bits of the sequence. Thus, masking with X=0 isolates

the strings 014 and 043 which define child node 0Xx, masking with X=1 isolates the strings 123,

173, 179, 152, 167 which defines the child node 1Xx, etc. Processing the Xxx root node

10 comprises distributing the thirteen strings into the child nodes 0Xx, 1Xx, etc. The child nodes

0Xx, 1Xx, etc. at Level 1 are next processed on the order 0Xx, 1Xx, etc. since $0 < 1 < \dots$ in

character value. Note that the characters are generally processed in the order 0, 1, 2, ..., 9 since 0

$< 1 < 2 < \dots$ in character value.

For the 0Xx node at level 1, the 0Xx mask is applied to the strings 014 and 043 to define

15 the next child nodes 01X and 04X, respectively, at Level 2. The 01X and 04X nodes are

processed in the sequential order of 01X and 04X since 0 is less than 4 in character value. Note

that the characters are always processed in the order 0, 1, 2, ..., 9. The 01X node at Level 2 is

processed, and since the 01X node contains only one string, the 01X node is a leaf node and the

string 014 is outputted to A(1). The 04X node at Level 2 is next processed and, since the 04X

20 node contains only one string, the 04X node is a leaf node and the string 043 is outputted to A(2).

For the 1Xx node at level 1, the 1Xx mask is applied to the strings 123, 152, 167, (173,

179) to define the next child nodes 12X, 15X, 16X, and 17X, respectively, at Level 2. The 12X, 15X, 16X, and 17X nodes are processed in the order 12X, 15X, 16X, and 17X, since the characters are always processed in the order 0, 1, 2, ..., 9 as explained *supra*. The 12X node at Level 2 is processed, and since the 12X node contains only one string, the 12X node is a leaf node and the string 123 is outputted to A(3). The 15X node at Level 2 is next processed and, since the 15X node contains only one string, the 15X node is a leaf node and the string 152 is outputted to A(4). The 16X node at Level 2 is next processed and, since the 16X node contains only one string, the 16X node is a leaf node and the string 167 is outputted to A(5). The 17X node at Level 2 is next processed such that the 17X mask is applied to the strings 173 and 179 to define the next child nodes 173 and 179 at Level 3, which are processed in the order of 173 and 179 since 3 is less than 9 in character value. The 173 node at Level 3 is next processed and, since the 173 node contains only one string, the 173 node is a leaf node and the string 173 is outputted to A(6). The 179 node at Level 3 is next processed and, since the 179 node contains only one string, the 179 node is a leaf node and the string 179 is outputted to A(7).

For the 2Xx node at level 1, since the 2Xx node contains only one string, the 2Xx node is a leaf node and the string 256 is outputted to A(8).

For the 5Xx node at level 1, the 5Xx mask is applied to the strings (512, 512), 561, 577, and 589 to define the next child nodes 51X, 56X, 57X, and 58X, respectively, at Level 2. The 51X, 56X, 57X, and 58X nodes are processed in the order 51X, 56X, 57X, and 58X, since the characters are always processed in the order 0, 1, 2, ..., 9 as explained *supra*. The 512X node at Level 2 is processed; since the node 51X does not include more than one unique string (i.e., 512

appears twice as duplicate strings), the 51X node at Level 2 is a leaf node and the duplicate strings 512 and 512 are respectively outputted to A(9) and (10). The 56X node at Level 2 is next processed and, since the 56X node contains only one string, the 56X node is a leaf node and the string 561 is outputted to A(11). The 57X node at Level 2 is next processed and, since the 57X node contains only one string, the 57X node is a leaf node and the string 577 is outputted to A(12). The 58X node at Level 2 is next processed and, since the 58X node contains only one string, the 58X node is a leaf node and the string 589 is outputted to A(13). Thus, the output array now contains the sorted strings in ascending order of value or pointers to the sorted values in ascending order of value, and the sorting has been completed.

Similar to the integer sort of FIG. 2, sorting the strings is essentially sorting the binary bits comprised by the strings subject to each character or byte of the string defining a unit of mask. Thus, the sorting algorithm is generally an algorithm for sorting sequences of bits whose interpretation conforms to the assumptions stated *supra*. No comparisons were made between the values of the strings to be sorted, which has the consequence of saving an enormous amount of processing time that would otherwise have been expended had such comparisons been made. The output array A(1), A(2), ..., A(13) contains sorted strings in ascending order of value as a consequence of the first assumption that for any two adjacent bits (or bytes) in the string to be sorted, the bit (or byte) to the left represents a larger magnitude effect on the value than the bit (or byte) to the right. If the alternative assumption had been operative (i.e., for any two adjacent bits (or bytes) in the string to be sorted, the bit (or byte) to the right represents a larger magnitude effect on the value than the bit (or byte) to the left), then the output array A(1), A(2), ..., A(8)

would contain the same strings as under the first assumption; however the sorted values in A(1), A(2), ..., A(8) would be in descending order of value.

Similar to the integer sort of FIG. 2, the preceding processes could be inverted and the sorted results would not change except possibly the ascending/descending aspect of the sorted strings in A(1), A(2), ..., (13). Under the inversion, the bytes 0, 1, 2, ..., 8, 9 would processed from right to left in the ordered sequence: 0, 1, 2, ..., 8, 9 (which is equivalent to processing the ordered sequence 9, 8, ..., 2, 1, 0 from left to right). As a result, the output array A(1), A(2), ..., A(8) would contain sorted strings in descending order of value is a consequence of the first assumption that for any two adjacent bits (or bytes) in the string to be sorted, the bit (or byte) to the left represents a larger magnitude effect on the value than the bit (or byte) to the right.

However under the inversion and if the alternative assumption had been operative (i.e., for any two adjacent bits (or bytes) in the value to be sorted, the bit (or byte) to the right represents a larger magnitude effect on the value than the bit (or byte) to the left), then the output array A(1), A(2), ..., A(8) would contain the sorted strings in ascending order of value.

As seen from the examples of FIGS. 2-4, the linked execution structure of the present invention includes nodes which are linked together in a manner that dictates a sequential order of execution of program code with respect to the nodes. Thus, the linked execution structure of the present invention may be viewed a program code execution space, and the nodes of the linked execution structure may be viewed as points in the program code execution space. Moreover, the sequential order of execution of the program code with respect to the nodes is a function of an ordering of masking results derived from a masking of the fields of the words to be sorted.

FIG. 5 is a flow chart for linear sorting under recursive execution, in accordance with embodiments of the present invention. The flow chart of FIG. 5 depicts the processes described *supra* in conjunction with FIGS. 2 and 4, and generally applies to sorting S sequences of binary bits irrespective of whether the sequences are interpreted as integers, floats, or strings. Steps 10-12 constitute initialization, and steps 13-20 are incorporated within a SORT module, routine, function, etc. which calls itself recursively in step 18 each time a new node is processed.

In step 10 of the initialization, the S sequences are stored in memory, S output areas A_1, A_2, \dots, A_S are set aside for storing the sorted sequences. S may be set to a minimum value such as, *inter alia*, 2, 3, etc. The upper limit to S is a function of memory usage requirements (e.g., see FIG. 16 and accompanying description) in conjunction with available memory in the computer system being utilized). The output areas A_1, A_2, \dots, A_S correspond to the output areas $A(1), A(2), \dots, A(S)$ described *supra* in conjunction with FIGS. 2 and 4. In addition an output index P and a field index Q are each initialized to zero. The output index P indexes the output array A_1, A_2, \dots, A_S . The field index Q indexes field of a sequence to be sorted, the field corresponding to the bits of the sequences that are masked and also corresponds to the levels of the linked execution structure.

In step 11 of the initialization, the root node E_0 is initialized to contain S elements associated with the S sequences. An element of a sequence is the sequence itself or a pointer to the sequence inasmuch as the nodes may contain sequences or pointers to sequences (e.g, linked lists) as explained *supra*.

In step 12 of the initialization, a current node E is set equal to the root node E_0 . The

current node E is the node that is currently being processed. Initially, the current node E is the root node E_0 that is first processed.

5 SORT begins at step 13, which determines whether more than one unique element is in the current node E being processed, is determining whether E is a leaf node. No more than one unique element is in E if E contains 1 or a plurality of identical elements, in which case E is a leaf node. If step 13 determines that there is no more than one unique element in E, then E is a leaf node and steps 14 and 15 are next executed. If step 13 determines that there is more than one unique element in E, then node E is not a leaf node and step 16 is next executed.

10 Step 14 outputs the elements of E in the A array; i.e., for each element in E, the output pointer P is incremented by 1 and the element is stored in A_p .

15 Step 15 determines whether the sort is complete by determining whether all nodes of the linked execution structure have been processed. Noting that SORT calls itself recursively in step 18 each time a new node is processed and that the recursed call of SORT processes only the values assigned to the new node, it is clear that all nodes have been processed when a normal exit from the first node processed by SORT (i.e., the root node) has occurred. Thus step 15 effectuates a normal exit from SORT. If said normal exit from SORT is an exit from processing the root node by SORT, then the sorting has ended. Otherwise, step 20 effectuates a return to execution of the previous copy of SORT that had been recursively executing. It should be noted that step 20 is not implemented by explicit program code, but instead by the automatic backward recursion to the previously executing version of SORT.

20

Step 16 is executed if E is not a leaf node. In step 16, the elements of E are distributed

into C child nodes: $E_0, E_1, \dots E_{C-1}$, ascendingly sequenced for processing purposes. An example of this is in FIG. 4, wherein if E represents the root node Xxx then the elements of E (i.e., the strings 014, 043, ..., 577, 561) are distributed into the 4 child nodes (i.e., $C=4$) of 0Xx, 1Xx, 2Xx, and 5Xx. The child nodes are ascendingly sequenced for processing, which means that the child nodes are processed in the sequence 0Xx, 1Xx, 2Xx, and 5Xx as explained *supra* in the discussion of FIG. 4.

Step 17 is next executed in which the field index Q (which is also the level index) is incremented by 1 to move the processing forward to the level containing the child nodes $E_0, E_1, \dots E_{C-1}$. Step 15 also initializes a child index I to 0. The child index points to the child node E_I ($I=1, 2, \dots, L$).

Steps 18-19 define a loop through the child nodes $E_1, E_2, \dots E_C$. Step 18 sets the node E to E_I and executes the SORT routine recursively for node E. Thus the child node E_I of the linked execution structure is a recursive instance of a point in the program code (i.e., SORT) execution space. When control returns (from the recursive call), the child index I is incremented by 1, followed in step 19 by a determination of whether the current child node E_I being processed is the last child to be processed (i.e., if $I=C$). If it is determined that $I \neq C$ then execution return to the beginning of the loop at step 18 for execution of the next child node. If it is determined that $I=C$ then all child nodes have been processed and step 20 is next executed. Step 20 effectuates a return to execution of the previous copy of SORT that had been recursively executing.

FIG. 6 is a flow chart for linear sorting under counter-controlled looping, in accordance with embodiments of the present invention. FIG. 6 effectuates the same sorting algorithm as FIG.

5, except that the procedure of FIG. 5 executes the nodes recursively, while the procedure of FIG. 6 executes the nodes iteratively through counter-controlled looping.

Step 31 provides initialization which may include substantially some or all of the processes executed in steps 10-12 of FIG. 5. The initializations in step 31 include storing the S sequences to be sorted, designating an output area for storing a sorted output array, initializing counters, etc. The number of sequences to be sorted (S) may be set to a minimum value such as, *inter alia*, 2, 3, etc. The upper limit to S is a function of memory usage requirements in conjunction with available memory in the computer system being utilized.

Step 32 manages traversal of the nodes of a linked execution structure, via counter-controlled looping. The order of traversal of the nodes are determined by the masking procedure described *supra*. The counter-controlled looping includes iterative execution of program code within nested loops. Step 32 controls the counters and the looping so as to process the nodes in the correct order; i.e., the order dictated by the sorting algorithm depicted in FIG. 5 and illustrated in the examples of 2 and 4. The counters track the nodes by tracking the paths through the linked execution structure, including tracking the level or depth where each node on each path is located. Each loop through the children of a level i node is an inner loop through nodes having a common ancestry at a level closer to the root. In FIG. 4, for example, an inner loop through the children 173 and 179 of node 17X at level 2 is inner with respect to an outer loop through nodes 12X, 15X, 16X, and 16X having the common ancestor of node 1XX at level 1. Thus, the inner and outer loops of the preceding example form a subset of the nested loops referred to *supra*.

Since the paths are complex and each path is unique, the node counters and associated child node counters may be dynamically generated as the processing occurs. Note that the recursive approach of FIG. 5 also accomplishes this tracking of nodes without the complex counter-controlled coding required in FIG. 6, because the tracking in FIG. 5 is accomplished automatically by the compiler through compilation of the recursive coding. Thus from a programming effort point of view, the node traversal bookkeeping is performed in FIG. 5 by program code generated by the compiler's implementation of recursive calling, whereas the node traversal bookkeeping is performed in FIG. 6 by program code employing counter-controlled looping explicitly written by a programmer. Using FIGS. 2, 4, and 5 as a guide, however, one of ordinary skill in the art of computer programming can readily develop the required program code (through counter-controlled looping) that processes the nodes in the same order as depicted in FIGS. 2, 4, and 5 so as to accomplish the sorting according to the same fundamental method depicted in FIGS. 2, 4, and 5.

Step 33 determines whether all nodes have been processed, by determining whether all counters have attained their terminal values. Step 33 of FIG. 6 corresponds to step 15 of FIG. 5. If all nodes have been processed then the procedure ends. If all nodes have not been processed then step 34 is next executed.

Step 34 establishes the next node to process, which is a function of the traversal sequence through the linked execution structure as described *supra*, and associated bookkeeping using counters, of step 32.

Step 35 determines whether the node being processed is empty (i.e., devoid of sequences

to be sorted or pointers thereto). If the node is determined to be empty then an empty-node indication is set in step 36 and the procedure loops back to step 32 where the node traversal management will resume, taking into account the fact that the empty node indication was set. If the node is not determined to be empty then step 37 is next executed. Note that steps 35 and 36
5 may be omitted if the coding is structured to process only non-empty nodes.

Step 37 determines whether the node being processed is a leaf node (i.e., whether the node being processed has no more than one unique sequence). Step 37 of FIG. 6 corresponds to step 13 of FIG. 5. If the node is determined to be a leaf node then step 38 stores the sequences (or pointers thereto) in the node in the next available positions in the sorted output array, and a
10 leaf-node indication is set in step 39 followed by a return to step 32 where the node traversal management will resume, taking into account the fact that a leaf node indication was set. If the node is not determined to be a leaf node then step 40 is next executed.

Step 40 establishes the child nodes of the node being processed. Step 40 of FIG. 6 corresponds to step 16 of FIG. 5

15 Step 41 sets a child nodes indication, followed by a return to step 32 where the node traversal management will resume, taking into account the fact that a child nodes indication was set.

Note that the counter-controlled looping is embodied in steps 32-41 through generating and managing the counters (step 32), establishing the next node to process (step 34), and
20 implementing program logic resulting from the decision blocks 33, 35, and 37.

Also note that although FIG. 6 expresses program logic to natural to counter-controlled

looping through the program code, while FIG. 5 expresses logic natural to recursive execution of the program code, the fundamental method of sorting of the present invention and the associated key steps thereof are essentially the same in FIGS. 5 and 6. Thus, the logic depicted in FIG. 6 is merely illustrative, and the counter-controlled looping embodiment may be implemented in any manner that would be apparent to an ordinary person in the art of computer programming who is familiar with the fundamental sorting algorithm described herein. As an example, the counter-controlled looping embodiment may be implemented in a manner that parallels the logic of FIG. 5 with the exceptions of: 1) the counter-controlled looping through the program code replaces the recursive execution of the program code; and 2) counters associated with the counter-controlled looping need to be programmatically tracked, updated, and tested.

FIGS. 7A, 7B, 7C, and 7D. (collectively "FIG. 7") comprise source code for linear sorting of integers under recursive execution and also for testing the execution time of the linear sort in comparison with Quicksort, in accordance with embodiments of the present invention. The source code of FIG. 7 includes a main program (i.e., void main), a function 'build' for randomly generating a starting array of integers to be sorted), a function 'linear sort' for performing the linear sort algorithm according to the present invention, and a function 'quicksort' for performing the Quicksort algorithm. The 'linear_sort' function in FIG. 7B will be next related to the flow chart of FIG. 5.

Code block 51 in 'linear_sort' corresponds to steps 13-15 and 20 in FIG. 5. Coding 52 within the code block 51 corresponds to step 20 of FIG. 5.

Code block 53 initializes the child array, and the count of the number of children in the

elements of the child array, to zero. Code block 53 is not explicitly represented in FIG. 5, but is important for understanding the sort time data shown in FIGS. 19-24 described *infra*.

Code block 54 corresponds to step 16 in FIG. 5.

5 Coding 55 corresponds to $I=I+1$ in step 18 of FIG. 5, which shifts the mask rightward and has the effect of moving to the next lower level on the linked execution structure.

Coding block 56 corresponds to the loop of steps 18-19 in FIG. 5. Note that linear_sort is recursively called in block 56 as is done in step 18 of FIG. 5.

10 FIGS. 8A, 8B, 8C, and 8D (collectively "FIG. 8") comprise source code for linear sorting of strings under recursive execution and also for testing the execution time of the linear sort, in comparison with Quicksort, in accordance with embodiments of the present invention. The coding in FIG. 8 is similar to the coding in FIG. 7. A distinction to be noted is that the coding block 60 in FIG. 8 is analogous to, but different from, the coding block 54 in FIG. 7. In particular, block 60 of FIG. 8 reflects that: a mask is not explicitly used but is implicitly simulated by processing a string to be sorted one byte at a time; and the string to be sorted may have a variable number of characters.

15 FIG. 9 illustrates a computer system 90 for sorting sequences of bits, in accordance with embodiments of the present invention., in accordance with embodiments of the present invention. The computer system 90 comprises a processor 91, an input device 92 coupled to the processor 91, an output device 93 coupled to the processor 91, and memory devices 94 and 95 each coupled to the processor 91. The input device 92 may be, *inter alia*, a keyboard, a mouse, etc. The output device 93 may be, *inter alia*, a printer, a plotter, a computer screen, a magnetic tape, a

removable hard disk, a floppy disk, etc. The memory devices 94 and 95 may be, *inter alia*, a hard disk, a dynamic random access memory (DRAM), a read-only memory (ROM), etc. The memory device 95 includes a computer code 97. The computer code 97 includes an algorithm for sorting sequences of bits in accordance with embodiments of the present invention. The processor 91 executes the computer code 97. The memory device 94 includes input data 96. The input data 96 includes input required by the computer code 97. The output device 93 displays output from the computer code 97. Either or both memory devices 94 and 95 (or one or more additional memory devices not shown in FIG. 9) may be used as a computer usable medium having a computer readable program code embodied therein, wherein the computer readable program code comprises the computer code 97.

While FIG. 9 shows the computer system 90 as a particular configuration of hardware and software, any configuration of hardware and software, as would be known to a person of ordinary skill in the art, may be utilized for the purposes stated *supra* in conjunction with the particular computer system 90 of FIG. 9. For example, the memory devices 94 and 95 may be portions of a single memory device rather than separate memory devices.

Timing Tests

FIGS. 10-24, comprise timing tests for the sort algorithm of the present invention, including a comparison with Quicksort execution timing data. FIGS. 10-15 relate to the sorting of integers, FIG. 16 relates to memory requirement for storage of data, FIGS. 17-18 relate to the sorting of strings, and FIGS. 19-24 relate to sorting integers as a function of mask width and

maximum value that can be sorted. The integers to be sorted in conjunction with FIGS. 10-15 and 19-24 were randomly generated from a uniform distribution. The timing tests associated with FIGS. 10-23 were performed using an Intel Pentium® III processor at 1133MHz, and 512M RAM.

5 FIG. 10 is a graph depicting the number of moves versus number of values sorted using a linear sort in contrast with Quicksort for sorting integers for a values range of 0-9,999,999. The linear sort was in accordance with embodiments of the present invention using the recursive sort of FIGS. 5 as described *supra*. For counting the moves, a counter was placed in the linear algorithm and in Quicksort at each point where a number is moved. Noting that 9,999,999
10 requires 24 bits to be stored, the linear sort was performed using mask widths $W=2, 3, 4, 6, 8, 12,$ and 14 with a corresponding number of levels $L=12, 8, 6, 4, 3, 2,$ and 2, respectively. For cases in which 24 is not an integral multiple of W , the mask width was truncated in the rightmost field corresponding to level L (i.e., at the level furthest from the root). For example at $W=14$, the mask widths at levels 1 and 2 were 14 and 10, respectively, for a total of 24 bits. FIG. 10 shows
15 that, with respect to moves for a values range of 0-9,999,999, Quicksort is more efficient than the linear algorithm for $W=2, 3,$ and 4, whereas the linear algorithm is more efficient than Quicksort for $W = 6, 8, 12,$ and 14.

 FIG. 11 is a graph depicting the number of compares/moves versus number of values sorted using a linear sort in contrast with Quicksort for sorting integers for a values range of 0-
20 9,999,999. For the linear sort, the number of compares/moves is the same as the number of moves depicted in FIG. 10 inasmuch as the linear sort does not “compare” to effectuate sorting.

For Quicksort, the number of compares/moves is a number of compares in addition to the number of moves depicted in FIG. 10. The linear sort was in accordance with embodiments of the present invention using the recursive sort of FIG. 5 as described *supra*. For counting the compares, a counter was placed in the linear algorithm and in Quicksort at each point where a number is compared or moved. Noting that 9,999,999 requires 24 bits to be stored, the linear sort was performed using mask widths $W=2, 3, 4, 6, 8, 12,$ and 14 with a corresponding number of levels $L=12, 8, 6, 4, 3, 2,$ and $2,$ respectively. For cases in which 24 is not an integral multiple of W , the mask width is truncated in the rightmost field corresponding to level L . For example at $W=14$, the mask widths at levels 1 and 2 were 14 and 10 , respectively, for a total of 24 bits. FIG. 11 shows that, with respect to compares/moves for a values range of $0-9,999,999$, the linear algorithm is more efficient than Quicksort for all values of W tested.

FIG. 12 is a graph depicting the number of moves versus number of values sorted using a linear sort in contrast with Quicksort for sorting integers for a values range of $0-9,999$. The linear sort was in accordance with embodiments of the present invention using the recursive sort of FIG. 5 as described *supra*. For counting the moves, a counter was placed in the linear algorithm and in Quicksort at each point where a number is moved. Noting that $9,999$ requires 14 bits to be stored, the linear sort was performed using mask widths $W=2, 3, 4, 6, 8, 10, 12, 14$ with a corresponding number of levels $L=7, 5, 4, 3, 2, 2, 2,$ and $1,$ respectively. For cases in which 14 is not an integral multiple of W , the mask width is truncated in the rightmost field corresponding to level L (i.e., in the cases of $W=3, 4, 6, 8, 10, 12$). FIG. 12 shows that, with respect to moves for a values range of $0-9,999$, Quicksort is more efficient than the linear

algorithm for $W=2, 3$, and 4 , whereas the linear algorithm is more efficient than Quicksort for $W = 6, 8, 10, 12$, and 14 .

FIG. 13 is a graph depicting the number of compares versus number of values sorted using a linear sort in contrast with Quicksort for sorting integers for a values range of $0-9,999$.

5 The linear sort was in accordance with embodiments of the present invention using the recursive sort of FIG. 5 as described *supra*. For counting the compares, a counter was placed in the linear algorithm and in Quicksort at each point where a number is compared. Noting that $9,999$ requires 14 bits to be stored, the linear sort was performed using mask widths $W=2, 3, 4, 6, 8, 10, 12, 14$ with a corresponding number of levels $L=7, 5, 4, 3, 2, 2, 2$, and 1 , respectively. For cases
10 in which 14 is not an integral multiple of W , the mask width is truncated in the rightmost field corresponding to level L (i.e., in the cases of $W=3, 4, 6, 8, 10, 12$). FIG. 13 shows that, with respect to compares for a values range of $0-9,999$, the linear algorithm is more efficient than Quicksort for all values of W tested. Of particular note is the difference in efficiency between the linear sort and Quicksort when the dataset contains a large number of duplicates (which
15 occurs when the range of numbers is $0-9,999$ since the number of values sorted is much greater than $9,999$). Because of the exponential growth of the number of comparisons required by the Quicksort, the test for sorting with multiple duplicates of values (range $0-9,999$), the test had to be stopped at $6,000,000$ numbers sorted.

FIG. 14 is a graph depicting the sort time in CPU cycles versus number of values sorted
20 using a linear sort in contrast with Quicksort for sorting integers for a values range of $0-9,999,999$. The linear sort was in accordance with embodiments of the present invention using

the recursive sort of FIG. 5 as described *supra*. Noting that 9,999,999 requires 24 bits to be stored, the linear sort was performed using mask widths W=2, 3, 4, 6, 8, 10, 12, and 14 with a corresponding number of levels L=12, 8, 6, 4, 3, 3, 2, and 2, respectively. For cases in which 24 is not an integral multiple of W, the mask width was truncated in the rightmost field corresponding to level L (i.e., at the level furthest from the root). For example at W=10, the mask widths at levels 1, 2, and 3 were 10, 10, and 4, respectively, for a total of 24 bits. As another example at W=14, the mask widths at levels 1 and 2 were 14 and 10, respectively, for a total of 24 bits. FIG. 14 shows that, with respect to sort time for a values range of 0-9,999,999, Quicksort is more efficient than the linear algorithm for W=2, 3, and 4, whereas the linear algorithm is more efficient than Quicksort for W = 6, 8, 10, 12, and 14.

FIG. 15 is a graph depicting the sort time in CPU cycles versus number of values sorted using a linear sort in contrast with Quicksort for sorting integers for a values range of 0-9,999. The linear sort was in accordance with embodiments of the present invention using the recursive sort of FIG. 5 as described *supra*. Noting that 9,999 requires 14 bits to be stored, the linear sort was performed using mask widths W=2, 3, 4, 6, 8, 10, 12, and 14 with a corresponding number of levels L=7, 5, 4, 3, 2, 2, 2, and 1, respectively. For cases in which 24 is not an integral multiple of W, the mask width was truncated in the rightmost field corresponding to level L (i.e., in the cases of W= 3, 4, 6, 8, 10, 12. FIG. 15 shows that, with respect to sort time for a values range of 0-9,999, the linear algorithm is more efficient than Quicksort for all values of W tested, which reflects the large number of compares for data having many duplicate values as discussed *supra* in conjunction with FIG. 13.

FIG. 16 is a graph depicting memory usage using a linear sort in contrast with Quicksort for sorting 1,000,000 fixed-length sequences of bits representing integers, in accordance with embodiments of the present invention using the recursive sort of FIG. 5 as described *supra*.

Quicksort is an in-place sort and therefore uses less memory than does the linear sort. The linear sort uses memory according to the following general formula, noting that this formula focuses only on the main memory drivers of the algorithm:

$$\text{MEM} = S * M_v + (M_c * 2^{W-1} * L)$$

wherein MEM is the number of bytes required by the linear sort, S is the number of sequences to be sorted, M_v is the size of the data structure (e.g., 12) required to hold each sequence being sorted, M_c is the size of the data structure (e.g., 8) required to hold a child sequence or pointer in the recursive linked execution structure, W is the width of the mask (≥ 1), and L is the number of levels of recursion. For some embodiments, $L = \text{ceiling}(M_v/W)$ as explained *supra*.

In FIG. 16, $M_v = 12$ and $M_c = 8$. The Quicksort curve in FIG. 16 is based on Quicksort using 4 bytes of memory per value to be sorted. The graphs stops at a mask width of 19 because the amount of memory consumed with the linear sort approaches unrealistic levels beyond that point. Thus, memory constraints serve as upper limit on the width of the mask that can be used for the linear sort.

FIG. 17 and 18 graphically depict the sort time in CPU cycles versus number of strings sorted for the linear sort and Quicksort, respectively. The linear sort was in accordance with

embodiments of the present invention using the recursive sort of FIG. 5 as described *supra*. The tests were conducted with simple strings. A file of over 1,000,000 strings was created by extracting text-only strings from such sources as public articles, the Bible, and various other sources. Each set of tests was run against strings ranging up to 20 characters in length (max_5 len=20) and then again against strings ranging up to 30 characters in length (max_len=30). A set of tests is defined as sorting a collection of 10,000 strings and repeating the sort with increasing numbers of strings in increments of 10,000. No sorting test was performed on more than 1,000,000 strings.

Quicksort is subject to chance regarding the value at the "pivot" points in the list of10 strings to be sorted. When unlucky, Quicksort is forced into much deeper levels of recursion (>200 levels). Unfortunately, this caused stack overflows and the tests abnormally terminated at 430,000 strings sorted by Quicksort. By reordering the list of strings, Quicksort could be made to complete additional selections, but the number of tests completed were sufficient to demonstrate the comparison of the linear sort versus the quicksort. FIGS. 17 and 18 shows that,15 with respect to sort time, the linear algorithm is more efficient than Quicksort by a factor in a range of about 30 to 200 if the number of strings sorted is at least about 100,000.

Another distinction between the linear sort and Quicksort is that in Quicksort the string comparisons define extra loops, which adds a multiplier A, resulting in the Quicksort execution time having a dependence of $A * S * \log S$ such that A is the average length of the string. The20 average length A of the string is accounted for in the linear sort algorithm as the number of levels L.

FIGS.17 and18 demonstrate that the linear sort far outperforms Quicksort for both max_ len=20 and max_len=30, and at all values of the number of strings sorted. A primary reason for the difference between the linear sort and Quicksort is that Quicksort suffers from a "levels of similarity" problem as the strings it is sorting become increasingly more similar. For example, to
5 differentiate between "barnacle" and "break", the string compare in the linear sort examines only the first 2 bytes. However, as Quicksort recurses and the strings become increasingly more similar (as with "barnacle" and "barney"), increasing numbers of bytes must be examined with each comparison. Combining the superlinear growth of comparisons in Quicksort with the increasing costs of each comparison produces an exponential growth effect for Quicksort.
10 Evidence of the effect of increasingly more costly comparisons in Quicksort can be understood by noting that the number of compares and moves made by the Quicksort are the same even though the maximum length of strings increases from 20 to 30. However, the number of clock cycles required to perform the same number of moves and comparisons in Quicksort increases (see FIG. 17) as the maximum length of strings increases from 20 to 30, because the depth of the
15 comparisons increases. FIG. 18 shows that the increase from 20 to 30 characters in the maximum length of strings affects the number of clock cycles for the linear sort, because the complexity of the linear sort is based on the size of the data to be sorted. The lack of smoothness in the Quicksort curves of FIG. 17 arises because of the sensitivity of Quicksort to the initial ordering of the data to be sorted, as explained *supra*.

20 FIGS. 19-24 is a graph depicting sort time using a linear sort, in contrast with Quicksort, for sorting integers as a function of mask width and maximum value that can be sorted, in

accordance with embodiments of the present invention. The values of S in FIGS. 19-24 are significantly smaller than the values of S used in FIGS. 10-15 and 17-18. The linear sort was in accordance with embodiments of the present invention using the recursive sort of FIG. 5 as described *supra*. In each of FIGS. 19-24, Time in units of CPU cycles is plotted versus MAX WIDTH and MOD_VAL, wherein MAX WIDTH (equivalent to W discussed *supra*) is the width of the mask, and wherein the integer values to be sorted were randomly generated from a uniform distribution between 0 and MOD_VAL-1. Also in each of FIGS. 19-24, MAX WIDTH = 13 is the rightmost array representing Quicksort and has nothing to do with a mask width. Letting S denote the number of integer values sorted in each test, S=2000 in FIGS. 19-20, S=1000 in FIGS. 21-22, and S=100 in FIGS. 23-24. FIGS. 19 and 20 represent the same tests and the scale of the Time direction differs in FIGS. 19 and 20. FIGS. 21 and 22 represent the same tests and the scale of the Time direction differs in FIGS. 21 and 22. FIGS. 23 and 24 represent the same tests and the scale of the Time direction differs in FIGS. 23 and 24. A difference between the tests of FIGS. 19-24 and the tests of FIGS. 10-16 is that much fewer values are sorted in FIGS. 19-24 than in FIGS. 10-16.

FIGS. 19-24 show a “saddle” shape effect in the three-dimensional Time shape for the linear sort. The saddle shape is characterized by: 1) for a fixed MOD_VAL the Time is relatively high at low values of MASK WIDTH and at high values of MASK WIDTH but is relatively small at intermediate values of MASK WIDTH; and 2) for a fixed MASK WIDTH, the Time increases as MOD_VAL increases.

Letting W denote MASK WIDTH, the effect of W on Time for a fixed MOD_VAL is as

follows. The Time is proportional to the product of the average time per node and the total number of nodes. The average time per node includes additive terms corresponding to the various blocks in FIG. 7B, and block 53 is an especially dominant block with respect to computation time. In particular, block 53 initializes memory in a time proportional to the maximum number of child nodes (2^W) per parent node. Let A represent the time effects in the blocks of FIG. 7B which are additive to the time ($\propto 2^W$) consumed by block 53. It is noted that 2^W increases monotonically and exponentially as W increases. However, the total number of nodes is proportional to N/W where N is the number of bits in each word to be sorted. It is noted that $1/W$ decreases monotonically as W increases. Thus the behavior of Time as a function of W depends on the competing effects of $(2^W + A)$ and $1/W$ in the expression $(2^W + A)/W$. This results in the saddle shape noted *supra* as W varies and MOD_VAL is held constant.

It is noted that the dispersion or standard deviation σ is inverse to the data density as measured by $S/(V_{MAX}-V_{MIN})$, wherein S denotes the number of values to be sorted, and V_{MAX} and V_{MIN} respectively denote the maximum and minimum values to be sorted. For FIGS 19-24, $V_{MIN} \geq 0$ and $V_{MAX} \leq MOD_VAL-1$. Thus, for a fixed data density of the S values, the Time is a saddle-shaped function of a width W of the mask. Although, FIGS. 19-24 pertain to the sorting of integers, the execution time of the linear sorting algorithm of the present invention for sorting sequences of bits is essentially independent of whether the sequences of bits are interpreted as integers or floating point numbers, and the execution time is even more efficient for string sorts than for integer sorts as explained *supra*. Therefore, generally for a fixed data density of S sequences of bits to be sorted, the sorting execution time is a saddle-shaped function of a width

W of the mask that is used in the implementation of the sorting algorithm.

At a fixed mask width W and a fixed number of values S to be sorted, increasing MOD_VAL increases the dispersion or standard deviation σ of the data to be sorted. Increasing σ increases the average number of nodes which need to be processed in the sorting procedure.

5 However, the Time increases as the average number of nodes needed to be processed increases. This results in the increase in Time as MOD_VAL increases while W is fixed. As to Quicksort, FIGS. 19-24 show that Time also increases as MOD_VAL increases for Quicksort.

A corollary to the preceding analyses is that for a fixed W, the standard deviation σ decreases (or the data density increases) as S increases, so that for a fixed W the sort execution
10 time may vary less than linearly with S (i.e., the sort execution time may vary as S^Y such that $Y < 1$).

FIGS. 19-24 show that for a given number S of values to be sorted, and for a given value of MOD_VAL, there are one or more values of W for which the linear sort Time is less than the Quicksort execution time. A practical consequence of this result is that for a given set of data to
15 be sorted, said data being characterized by a dispersion or standard deviation, one can choose a mask width that minimizes the Time and there is one or more values of W for which the linear sort Time is less than the Quicksort execution time.

Although FIGS. 19-24 shows timing tests data for sorting integers, the ability to choose a mask resulting in the linear sort of the present invention executing in less time than a sort using
20 Quicksort also applies to the sorting of floating point numbers since the linear sort algorithm is essentially the same for sorting integers and sorting floating point numbers. Additionally, the

ability to choose a mask resulting in the linear sort executing in less time than a sort using Quicksort also applies to the sorting of character strings inasmuch as FIGS. 14-15 and 17-18 demonstrate that the sorting speed advantage of the linear sort relative to Quicksort is greater for the sorting of strings than for the sorting of integers. It should be recalled that the mask used for the sorting of character strings has a width equal to a byte representing a character of the string.

While embodiments of the present invention have been described herein for purposes of illustration, many modifications and changes will become apparent to those skilled in the art. Accordingly, the appended claims are intended to encompass all such modifications and changes as fall within the true spirit and scope of this invention.